

STOCHASTIC-BASED ADAPTIVE CONTROL VIBRATION CONTROL FOR MACE II

Daniel P. Scharf* and David C. Hyland†

Department of Aerospace Engineering

University of Michigan, Ann Arbor, MI 48109-2140

Abstract

The Middeck Active Control Experiment (MACE) is being reflown as MACE II. Its purpose is to gain on-orbit experience with a variety of adaptive controllers. The University of Michigan contribution to the effort is Fixed-Structure Adaptive-Improvement (FSAI) control. FSAI is a framework for online control optimization which is designed to augment a baseline, low-gain robust controller so as to improve its performance: FSAI is adaptive control aimed at stable plants. To obtain Bounded-Input Bounded-State (BIBS) stability, FSAI proceeds by generating controllers via an optimization algorithm, selectively implementing them using an optional, online Nyquist check, evaluating the performance online, opening the loop if at any time the performance exceeds the open loop (plant+baseline) performance, and then allowing the stable plant to settle so as to return the plant state to the same neighborhood of the origin (thus ensuring switched system stability). Implemented on the MACE II hardware on the ground, in 6 separate runs FSAI control generated over 32,000 controllers in 3.4 hours with average improvements of the baseline designs of 37% and without the output saturating. A straightforward, automatic, online frequency-domain identification algorithm was also developed as part of the online Nyquist check.

Introduction

An inherent problem in spacecraft control design is that development and testing occurs in a 1-g environment whereas the spacecraft operates in 0-g. Even with gravity offload systems gravity still: loads the structure causing changes in the structure's stiffness; loads joints, thereby reducing their nonlinear behavior; and deforms the structure, resulting in the structure's dynamics being measured with respect to a non-0-g equilibrium.⁴ In addition, the dynamics of the offload system must now be considered. For one flight experiment the frequencies of the first 10 modes shifted in going from 1-g to 0-g from 0.4 to 12% with 5 of the modes shifting more than 3%.³

In other cases, it is not possible to test the spacecraft as a whole. Finite element method (FEM)

models are developed for spacecraft components based on technical specifications and material properties, and when possible, the models are refined using test data. The separate models are then combined to form a FEM model for the entire spacecraft. A drawback to this procedure is that the overall FEM model may be too inaccurate for use in high-performance control design. For NASA's Galileo spacecraft, 1-g data only gave the modal frequencies to within 20% and the mode shapes to within 50%.¹

Performance requirements often dictate actively controlling flexible modes of a spacecraft as opposed to limiting control bandwidth to avoid control-structure interaction. In trying to control these uncertain, flexible modes two approaches have been considered: robust control and adaptive control. In the first case, a bound on the uncertainty of the structural dynamics is developed and the control is forced to be stable over the entire range of possible dynamics. By their nature robust controllers sacrifice performance for a greater tolerance to mis-modeled/unknown dynamics.

Adaptive controllers have been examined in the hopes that this performance may be recovered. As the adaptive controller is "tuning" itself to the actual, 0-g dynamics, robustness is no longer required; in some sense the adaptive controller learns the actual dynamics and, therefore, is not required to be stable over a possibly conservatively wide range of dynamics.

The Middeck Active Control Experiment (MACE), which flew on STS-67 in 1995, developed and validated tools for predicting 0-g dynamics and control performance from 1-g data. Specifically, MACE considered the pointing of a payload at the end of a flexible appendage subjected to stochastic disturbances. The approach taken³ was to construct an accurate FEM model from 1-g data and determine the range of parametric uncertainty in the 1-g model. Next the gravity offload system and gravity itself were "turned off" in the model producing a 0-g FEM model. Finally, the 1-g uncertainty estimates were projected into estimates of the 0-g parametric uncertainty. Using a variety of robust control synthesis techniques, controllers were designed based on the 0-g FEM model incorporating the 0-g uncertainty predictions.⁸ MACE also performed on-orbit identification with the results being sent to the ground where con-

*dscharf@umich.edu

†Professor and Chair, dhyland@engin.umich.edu

Copyright © 2001 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved.

trollers were redesigned. The redesigned controllers were uploaded and also run.

MACE was a successful mission demonstrating the accuracy and efficacy of the 0-g modeling and robust control techniques used. A lesson learned from MACE was that, while the on-orbit performance of low-to-medium and even some higher gain controllers could be reliably predicted, the highest gain designs required on-orbit data and redesign to be stably implemented (an unexpectedly high frequency tether mode destabilized the ground-designed controllers). Also noted was that larger, more complex spacecraft structures may not be amenable to the extensive preflight testing used to develop the accurate FEM model necessary for successful 0-g predictions. In these cases, coarser models could be used to obtain a modest improvement in pointing performance and on-orbit identification coupled with ground-based controller redesign would be used to obtain the performance requirements.¹⁵

MACE II is a flight experiment using a refurbished MACE test article to examine the capabilities of adaptive controllers for pointing control and line-of-sight (LOS) jitter reduction. One of the goals of MACE II is to develop adaptive algorithms that greatly reduce the extensive modeling, noted previously, and on-orbit testing necessary to successfully implement robust controllers for high-performance missions.¹⁶ The Air Force Research Laboratory (AFRL), which is responsible for MACE II, has invited teams from industry and academia to provide adaptive controllers. Reported herein is the adaptive controller developed by the University of Michigan (UM) team and the experimental results of this controller in preflight, ground tests.

The MACE II test article will be briefly described, followed by a statement of the control problem addressed by the UM team. The adaptive controller developed, referred to as Fixed-Structure Adaptive-Improvement (FSAI) control, will be outlined and some theoretical properties discussed. Finally, experimental results will be presented.

MACE II Hardware Description

The MACE configuration was designed to represent an actively controlled, high-payload-mass-fraction spacecraft. The test article consists of a flexible bus to which two payloads, a reaction wheel assembly (RWA) and other actuators and sensors are mounted (see Figure 1). Each payload is mounted to the structure using a two-axis gimbal that provides pointing capability. Instrumentation includes angle encoders on each gimbal axis, a three-axis rate gyro platform under the RWA and a two axis rate gyro

platform mounted in the primary payload. The bus is composed of circular cross-section struts connected by aluminum nodes. The structure is supported for ground tests by a pneumatic/electric suspension system.

Because the suspension cables, the bus and gravity vectors are all in the same plane, the structural dynamics decouple into vertical (about the Z-axis) and horizontal (about the X and Y axes) dynamics. The control is implemented using a 33 MHz processor operating at a 500 Hz sampling rate. Approximately 0.8 megabytes are available for the controller executable file and any adaptive controller memory requirements.

The UM team considered a SISO vibration control problem where an unmeasurable, broadband (white noise filtered through a Butterworth filter with a corner at 50 Hz) disturbance is injected into the secondary payload's Z torque actuator (SGZt), the primary payload's torque actuator (PGZt) is used for control and the primary payload's Z gimbal rate gyro (PGRGZ) is the output measurement. The PGZt and PGRGZ are colocated, but anti-aliasing filters and computational delay introduce an approximately 5 time step delay rendering the PGZt-to-PGRGZ loop non-minimum phase.

The performance metric for the MACE II flight experiment is the square root of the integral from 0.5 to 50 Hz of the power spectral density (PSD) of the integrated PGRGZ signal (the output). This metric is the bandlimited RMS of the Z-axis angle of the primary payload (the integration to obtain angle from rate is done in the frequency domain).

A modified metric was used for the UM team's controller development. Evaluating the performance metric described above takes 114 seconds, and obtaining a good approximation takes around 20 seconds. FSAI control, as will be described, is essentially an online optimization framework—the more controllers that can be tried during the experimental run, the more improvement will be observed. Using the above metric allows approximately 300 controllers to be evaluated during an experimental run. Instead, if the bandlimited RMS of the Z-axis *angle rate* is used, upwards of 5000 controllers can be evaluated. Integrating the angle rate weights the lower frequencies with respect to the higher frequencies, resulting in the need to take longer data streams to capture performance relevant information. Using the rate instead allows performance relevant information to be obtained in less than a few seconds.

Simulations were run with the actual MACE II metric (the bandlimited Z-axis angle RMS), but due to programming delays, ground data was not collected




Figure 1: MACE II Test Article Schematic

using the angle metric. The ground results to be presented are for the rate metric.

Control Problem

The control problem considered for MACE II may then be stated as follows: develop an adaptive algorithm to regulate a non-minimum phase, asymptotically stable, unknown, LTI plant subjected to an unmeasurable, broadband disturbance. While a structure on orbit is not open loop asymptotically stable due to rigid body modes (RBMs), it is assumed a low-bandwidth controller stabilizes these modes. This is equivalent to assuming a baseline, stabilizing robust controller is running which provides some performance, but, as mentioned in the MACE lessons learned, it is of low-to-medium gain. The open loop is then redefined to be the open loop plant plus baseline controller in feedback. The goal is to augment the baseline, robust controller, recovering the performance traded for robust stability guarantees.

Review of Adaptive Vibration Control

As the disturbance is broadband and unmeasurable schemes aimed at tonal cancellations and feedforward controllers are not considered. Robust Adaptive Control (RAC) is the field of adaptive control which treats disturbances and/or unmodelled dynamics. Direct RAC schemes, for example, modified Model Reference Adaptive Control (MRAC), require the plant to be minimum phase.¹¹ Indirect RAC schemes, such as the Adaptive Pole Placement algorithms collected in Ioannou and Sun,¹¹ are applicable. A drawback is that bounds on performance are not practically computable. For instance, it can be shown (Ioannou and Sun, Theorem 9.5.2) that the output is bounded as $\int_0^t y^2(\tau) d\tau \leq c(d_0^2 + f_0) + c$ where f_0 is a constant depending on the method by which the adaptive controller is made robust (e. g. dead zone), d_0 is a bound on the disturbance and c is a constant which can be shown to exist, but as far as the authors

are aware, cannot be computed.

However, RAC addresses a different problem: the plant is assumed to be open loop unstable. Assuming that the plant is open loop asymptotically stable (possibly due to a baseline, robust controller in feedback) should allow more to be said about performance.

If the plant is open loop stable, the purpose of adaptive control changes from stabilization to *performance enhancement*. This emphasis is discussed in Wang, Mareels and Moore²² and Tay and Moore.¹⁹ Adaptive control aimed at performance enhancement will be referred to as Performance Adaptive Control (PAC). Wang et al., Tay and Moore and others (Venu-gopal and Bernstein, Trulsson and Ljung) propose and analyze what may be called direct PAC schemes: a performance criterion is minimized online through direct modification of the control parameters. These schemes generally require the assumption of bounded signals to show convergence of the performance metric to the minimum or are of a local nature—if the “true” plant is sufficiently close to the “nominal” plant, the metric will converge (not necessarily to the optimum).

Indirect PAC schemes, as exemplified by Bayard, Yam and Mettler,² iterate and generally couple, identification and control design. While referred to as “adaptive,” these methods embody the principle of successive controller refinement. The final design of the iteration, however, remains a robust design by necessity (part of the identification process provides error bounds for robust design).

For direct PAC schemes, a method for ensuring the boundedness of signals is needed, and for indirect PAC schemes, there is still the possibility of recovering the performance lost in robust designs.

FSAI Algorithm

Building upon the idea of successive controller refinement, Fixed-Structure Adaptive-Improvement (FSAI) control separates performance optimization

and stability concerns, providing a framework for online controller improvement. It at best improves the baseline controller performance because the performance metric problem may be multimodal (contain multiple local minima) and be high-dimensional (many controller parameters) and such optimization problems are presently an open problem.¹⁰ Stochastic algorithms do guarantee convergence with probability one, but can take a great deal of time to converge.

FSAI control has two parts. The Optimization Module (OM) implements an optimization routine of the designers choice—there are no restrictions. Further, the parameterization of the controller can also be chosen freely. Bounded-Input Bounded-State (BIBS) stability is ensured by the Stability Evaluation Module (SEM).

The FSAI control algorithm is now described and then its BIBS stability is discussed.

The assumptions on the plant are as follows:

- A1) The plant model is linear, time-invariant (LTI) and characterized by an unknown parameter vector $\theta \in \mathbb{R}^{n_\theta}$.
- A2) $\theta \in \Theta$, where Θ is known and compact.
- A3) For all $\theta \in \Theta$ the plant model is asymptotically stable.
- A4) The plant model is strictly proper.
- A5) The plant model is subjected to a bounded input disturbance and the bound is known.

Note that the plant is assumed to be parametrically uncertain, and the true parameter vector belongs to a known, compact set. There are a number of identification algorithms which provide such a compact set, for example, the ellipsoidal sets of Kosut, Lau and Boyd.¹³

A controller parameterization is assumed to exist so that the controller transfer function may be written as $C(s, \mathbf{p})$, where \mathbf{p} is the controller parameter vector. Figure 2 shows the FSAI controller is a block diagram. A plant is subjected to an input disturbance d and sensor noise v with the FSAI controller in feedback. u is the control input and y the plant output.

The performance metric is $J = J(\mathbf{p})$. In analogy to the output constraint set of Gilbert and Tan,⁷ which is a set Y such that if the output remained in the set, then the state and control are considered acceptable (i. e. the system under consideration was in a safe operating condition), we will enforce the principle that the closed loop performance should not be worse than the open loop performance. It is assumed that the performance metric is monotonic non-decreasing in time, becomes unbounded as the output becomes unbounded and only requires a finite time, T_J , to evaluate from output measurements. For example, the

performance metric used in the MACE II experimental runs to be presented was $J = \sum_{k=1}^N y_k^2$, where $N = 600$ time steps and y_k is the measurement at time step k .

The FSAI algorithm is as follows (refer to Figure 2):

1. The OM evaluates the open loop performance, J_{ol} , and passes it to the Performance Enforcement Block (not shown in figure).
2. The OM generates a controller \mathbf{p} .
3. \mathbf{p} is passed to an optional, online Nyquist check.
4. If \mathbf{p} passes the check, it is implemented.
5. The controller runs until the performance is evaluated or $J(\mathbf{p}) > J_{ol}$.
6. If the controller's performance is fully evaluated, it is passed to the OM, goto step 2.
7. If the controller's performance becomes greater than J_{ol} , the feedback loop is opened (via the switch in the figure controlled by the Performance Enforcement Block), and the plant is allowed to settle a predetermined amount of time. After settling, goto step 2.

This process continues until either the performance metric has been calculated a prespecified number of times or satisfactory performance is achieved.

Though it is not theoretically necessary for BIBS stability, an online Nyquist check makes the FSAI controller much more efficient in terms of performance improvement versus time, and greatly improves the time history of the output as fewer destabilizing controllers are implemented. The Nyquist check uses a frequency response function (FRF) of the plant that is identified online automatically (the ID algorithm is discussed in the MACE II Ground Results section). Since $C(s, \mathbf{p})$ is known, it and the plant FRF are used to generate an estimate of the loop gain in the frequency domain. A simplified Nyquist check is used: the loop gain is not allowed to enter a "lock-out" region about the -1 point. This type of check simplifies the logic and eliminates conditionally stable and unstable controllers.




Figure 3: Lock-out Region for Nyquist Criterion

We note that it is theoretically possible to determine stability from a finite number of points in the frequency domain,¹² but now the frequency response of the characteristic polynomial must be considered (this would require fitting a transfer function to the plant FRF to obtain the numerator and denominator

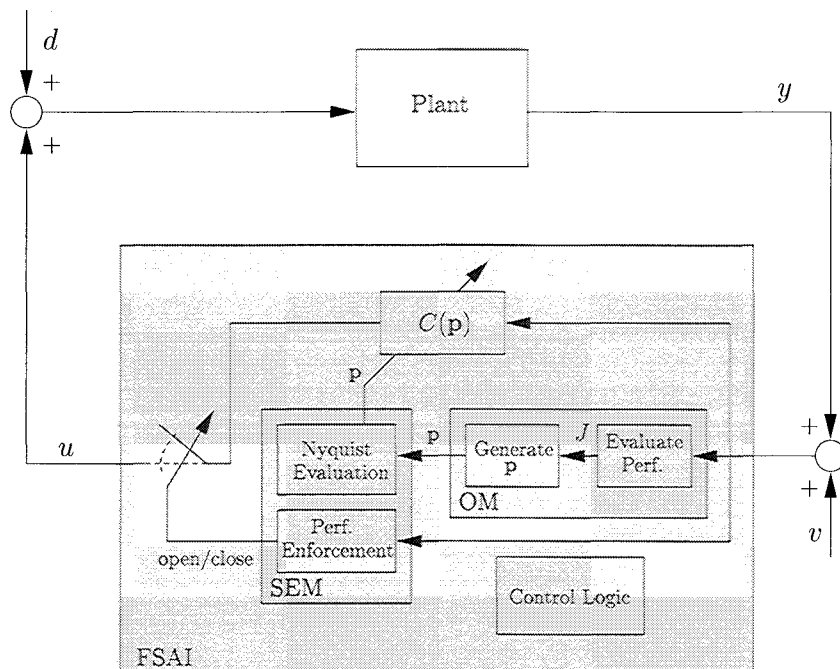


Figure 2: FSAI Architecture

polynomials). More importantly, when a disturbance and noise are present, the identified model would require error bounds to ensure stability, and thus the finite information test becomes a robust stability test, which, in turn, removes the possibility of recovering performance.

BIBS Stability

The FSAI algorithm just presented is the algorithm as it was implemented on MACE II. It does not have the theoretical guarantee of BIBS stability. To obtain this guarantee without adding more assumptions, the plant must be allowed to settle after each and every controller implementation whether or not the performance evaluation ran to completion or was terminated early due to $J(\mathbf{p}) > J_{ol}$. The reason is that the FSAI algorithm is a switched system.

From Liberzon and Morse,¹⁴ a switched system is a family of state matrices $\{A_q : q \in Q\}$, a rule $\sigma(t) : [0, \infty) \rightarrow Q$ that is piecewise constant in time and the state equation $\dot{x} = A_{\sigma(t)}x$. Even if all the A_{σ} (referred to as subsystems) are asymptotically stable, it is possible for the system to have an unbounded response. That is, even if FSAI control only implements stabilizing controllers, it is still possible that the state could grow without bound. As no assumptions of observability have been made, it is possible for the performance to remain bounded even as the state becomes unbounded.

Results ensuring BIBS stability for switched systems are generally of the form: if one switches in-

frequently enough and if the ratio of stable subsystem activation to unstable subsystem activation is large enough, then switched system is exponentially stable.²⁵ However, the “enough” quantities depend on a Jordan-form-based bound of the matrix exponential of the subsystems and the ratio of the spectral abscissa of the most unstable subsystem to that of the least stable subsystem, respectively, both of which can be quite conservative. Further, as $J(\mathbf{p}) < J_{ol}$ implies nothing about the location of the eigenvalues of the closed loop state matrix (unstable systems can have arbitrarily good performance over finite time scales) and the Nyquist check cannot guarantee stability (due to identification error and/or finite information), each controller implemented must be assumed to be destabilizing, and hence the ratio of stable subsystem activation to unstable subsystem activation must be taken as the ratio of time the plant is allowed to settle to the time a controller is implemented. All of these factors lead to excessively long settling times.

A different approach is taken for FSAI. Instead of trying to obtain exponential stability, BIBS stability is obtained by returning the plant state to the same neighborhood of the origin (in phase space) between each controller implementation. This is done by using a bound on the matrix exponential on a set of matrices.

Recall $\theta \in \Theta$, where Θ is compact. The state of all plants in the uncertainty set can be bounded if a bound on $\|e^{A\theta t}\|$ can be found for all state matrices A resulting from $\theta \in \Theta$. It is possible to extend

current bounds on the matrix exponential to sets of matrices in many cases. However, when considering a set of matrices which contain defective matrices (matrices with at least one Jordan block of order greater than one), only the Schur Decomposition-based bound can be extended, and it can be quite conservative.²¹ A new bound based on a combination of the Jordan- and Schur-based bounds has been developed which is less conservative in many cases and it is referred to as the Hybrid bound. Using the Hybrid bound and the history of the control used during a controller performance evaluation, it is possible to calculate a bound on the settling time to return the plant state to a neighborhood of the origin. The algorithm only involves bisection of $(n - 1)$ th order polynomial, and so can be done relatively inexpensively online.

The derivation of the Hybrid bound and the settling algorithm is involved and is not reproduced here. Details can be found in Scharf.¹⁸

With the bound on settling time, each controller is implemented when the plant state is in the same, known neighborhood of the origin. Using the controller saturation value and T_J , the maximum length of time any controller is implemented, the maximum magnitude of the plant state of any plant in the uncertainty set can be calculated. Then, using the Hybrid bound, the maximum transient in settling from this maximum state magnitude can also be found. And so, the FSAI control algorithm is BIBS stable.

However, the settling times are still found to be too conservative for practical implementation. One area for further inquiry is what can be said if observability is assumed for all plants in the uncertainty set. For MACE II, the pragmatic approach of finding a suitably long settling time through experiment was taken (a settling time of 0.2 seconds was selected). Further, the plant was allowed to settle only if $J > J_{ol}$, otherwise the next controller generated by the OM was implemented without pause. This modified algorithm was the one enumerated earlier.

Optimization Module

The OM generates the controllers for online improvement. With BIBS stability provided (assuming that variant of the algorithm is used) by the Stability Evaluation Module (SEM), any optimization algorithm and controller parameterization can be used to attempt to optimize, online, the performance metric.

The performance metric, as discussed in the MACE II Hardware section, is the PGRGZ RMS from 0.5 to 50 Hz. This metric is not monotonic non-decreasing in time, so for online optimization a metric which approximates the bandlimited RMS, $J(\mathbf{p}) = \sum_{k=1}^{600} y_k^2(\mathbf{p})$, is used. In the ground results

to be presented, y_k is not bandpass filtered, so the online metric is truly only an approximation, but it is easy to calculate and is shown in the next section that improvement can be achieved with it.

The parameterization $C(s, \mathbf{p})$ used by the UM team is based upon classical design concepts and is referred to as General Modal Control (GMC): a number of basic filter blocks are combined to form a “general” form to control one mode, and successive general forms are connected in parallel to control successive lightly damped modes. This approach is similar to the generalized structural filtering concept of Wie and Byun.²³ In the form chosen \mathbf{p} has dimension 23, that is, there are 23 controller parameters to be selected, and $C(s, \mathbf{p})$ is 13th order.

Optimizing a 23rd dimensional performance metric, possibly having local minima, is a difficult problem. Many deterministic optimization algorithms become impractical for dimensions greater than 10.^{10,20} Stochastic algorithms were used in the OM as they can be applied to a wide number of problems with little modification and have been used on problems with dimensions greater than 100.²⁰ However, their convergence can be quite slow.

Instead of trying to solve the 23rd dimensional optimization problem online, a solution of which is not assured offline, the Super-Heuristic method of Ho⁹ was adopted with modification. Super-Heuristics (SH) is based upon the simple idea that a control design “...born out of human experience, insight, and ingenuity can often be good enough but is seldom optimal...,” and as such, there is a good chance that perturbing the heuristic design will lead to a better design. Note the lack of optimality may come from using classical design techniques or even using optimal design methods with truncation or with an uncertain plant model.

Ho⁹ uses the following example as justification. Given a design let ϵ be the probability that perturbing the design a small amount leads to a better design. Then the probability of obtaining a better design in N perturbations is $1 - (1 - \epsilon)^N$. For $\epsilon = 0.05$ and $N = 20$, the chance of getting a better design is greater than 50%.

We take the idea of Super-Heuristics one step further—let the perturbations be based upon the controller design methodology. The GMC controller controller used on MACE II is essentially two individual “mode controllers” in parallel, each mode controller consisting of a bandpass filter and two notches (which increase the classical stability margins). The design proceeds by designing the bandpass filter of the first mode controller to attenuate the response of the first structural mode. The two notches are then designed

to increase the classical stability margins *and then* the bandpass filter is redesigned to take advantage of the increased margins. The bandpass-notch-bandpass process is repeated for the second mode controller applied to the second structural mode. Then the entire process is repeated starting again with the first bandpass filter. Refinement continues in this manner until performance requirements are satisfied or stability limitations prohibit increasing the gain further. Also, a lead-lag filter is used to shape the overall phase and a proportional feedback gain is included to improve damping.

The GMC controller designed was compared to a truncated linear quadratic Gaussian (TLQG) design on a 30th order auto-regressive, moving-average (ARMA) model of the MACE II PGZt to PGRGZ loop obtained from experimental data. The TLQG design was synthesized using the disturbance and noise magnitudes of the MACE II experiment with the LQ weights being chosen so that the TLQG design had the same control RMS as the GMC design. That is, both the GMC and TLQG designs use the same control effort. The TLQG design was obtained through a balanced truncation of the LQG design to 13th order. The performance metric for comparison is simply the steady-state output variance. The value of the metrics, the classical stability margins and the value of the metrics after Super-Heuristics (SH) and Modified Super-Heuristics (MSH) have been applied to GMC and TLQG designs are shown in Table 1. Note that MSH cannot be applied to the TLQG design as no rationale was apparent for choosing subvectors of \mathbf{p} to perturb. Also, 2σ error bounds are shown in the table. The discrete Lyapunov equation for calculating the steady-state output variances was ill-conditioned and so the metrics were calculated through time averages. While SH does better in this example than MSH in improving the GMC design, MSH was found to be slightly more promising for implementation.

Table 1: Comparison of Initial and Improved GMC and TLQG Designs

| Property | GMC | TLQG |
|-------------------------------------|-------------|-------------|
| Initial Cost (2σ) | 26.70(1.12) | 24.24(1.17) |
| Initial Gain/Phase Margins, dB/deg. | 6.1/67 | 6.5/128 |
| SH | | |
| Final Cost (2σ) | 21.02(0.22) | 24.19(0.31) |
| MSH | | |
| Final Cost (2σ) | 21.38(0.36) | N/A |

While initially better, the TLQG design is very near a local minimum of the performance metric. This example shows that improvement obtained in the ex-

perimental results of the next section is not simply due to having a very poor initial design; the GMC controller compares well to a truncated LQG design. It also compares favorably to an LQG design used on the ground experiments for the first MACE mission.⁸ Grocott et al.⁸ do not state the order of this LQG controller, but the original MACE software could run an 80th order compensator.

MACE II Ground Results

The final version of the code delivered to the Air Force Research Laboratory (AFRL) for the flight experiment was run on six cases. Five PGZt to PGRGZ cases were run, four with the disturbance injected into the SGZt and one with the disturbance injected into the PGZt. A case for the PGXt-to-PGRGX loop was also run with the disturbance injected into the SGXt. The Z loop GMC controllers were not redesigned for the different disturbance sources, and for all loops the performance improvement was comparable. A disturbance injected into the a primary gimbal will be referred to as D1 and into a secondary gimbal as D2.

The experimental runs take approximately 55 minutes. The automatic online ID, discussed subsequently, takes 17 minutes, the control improvement runs for 35 minutes on average, and initialization and data collection requirements take 3 minutes. In the 35 minutes of control improvement, 5400 controllers are generated by the Modified Super-Heuristics (MSH) algorithm and are selectively implemented by the online Nyquist check.

Online Identification

The frequency response function (FRF or transfer function) of the plant is estimated using the well known cross-spectral power method (Franklin et al.,⁵ pg. 492). The algorithm for calculating the power spectra is taken from Press et al.¹⁷ In addition, we also 1) calculate the coherence online, and if it is not acceptable the frequency range wherein the coherence is sub-par is broken into two equal sub-regions and the identification is repeated in successive sub-regions until the coherence tolerance is met or the frequency ranges become smaller than another tolerance, and 2) perform a pre-ID scan to maximize the signal-to-noise ratio (SNR) of the identification signal. This scan is done by repeatedly increasing the magnitude of the ID signal until either the input or output saturate. The ID signal magnitude is then decreased by a safety factor to ensure saturation does not occur during data collection for identification.

Coherence checks and frequency “zooming,” as the breaking down of the frequency range is sometimes called, are tasks commonly carried out during an iter-

ative, designer-in-the-loop modeling effort. We have automated them in a simple manner.

Bayard et al.¹ develop a comprehensive automatic, online ID algorithm aimed at the specific problems of on-orbit, large space structure identification. The identification inputs can be broadband or narrowband stochastic, or sinusoidal. They also use stochastic spectral estimation using Welch's Periodogram Method. Bayard et al. also mention that if the actuator power is constrained, the frequency range can be broken into several, overlapping narrowband (via their narrowband stochastic ID signals) ranges. The estimates in each individual ranges can then be combined into an FRF estimate over the entire frequency range.

Yam, Bayard and Scheid²⁴ use the Schroeder-Phased Sinusoidal Input Design for their ID signal. Discrete Fourier Transforms (DFTs, used in Welch's Method) are not exact; the value at a certain bin frequency is actually the average of not only that bin, but also of the surrounding bins. This phenomenon is known as leakage. One of the implications of the averaging function is that if a signal only contains spectral content at the bin frequencies of a DFT, then no leakage occurs. If there are no disturbances, or if the ID signal is large compared to them, and if the ID signal is a sum of sinusoids at the bin frequencies, then the DFT will be exact. That is, there is no discrepancy between the continuous Fourier Transform and the DFT.

One drawback to the "sum-of-bin-frequency-sinusoids" (SBFS) approach is that the ID signal has very large peaks in the time domain. The Schroeder Design phases the bin-frequency sinusoids so that this peaking is minimized. We do not Schroeder phase the sinusoids in our SBFS ID signal as we were unaware of this result until after the FSAI software had been submitted to the AFRL. The FSAI identification algorithm would almost certainly benefit from it.

The automated ID algorithm for MACE II combines the frequency sub-range concept mentioned in Bayard with the SBFS concept mentioned, among other places, in Yam et al. and Press et al. The PSD of a SBFS signal looks like a comb, and so we refer to the frequency sub-ranges as combs. Given an initial number of combs, N_c , the ID algorithm breaks the entire frequency range into N_c equally sized combs. The SBFS signals for each comb are generated using nominal magnitude scalings. Then for each comb, the ID signal is increased until saturation of the input or output. This process is the pre-ID scan mentioned earlier.

After the pre-ID scan, the ID signal for each

comb is repeatedly applied to the plant a user-specified number of times and the output is averaged in the time domain.

Parameters for the MACE II implementation are 4 combs, 1024 point DFTs (0.49 Hz spacing) and the final FRF is the result of the average of two FRFs, each FRF based on data that has been averaged 50 times in the time domain. With these parameters the ID data collection takes 13.7 minutes, the pre-scan takes on average approximately 0.9 minutes and the computation of the DFTs and periodograms takes the remaining 2.4 minutes.

FSAI Results

A time history of the PGRGZ signal for a FSAI experimental run is shown in Figure 4. First, during the "Init" stage, start-up transients are allowed to die out and all memory allocation and initialization takes place. Then, open loop (OL) performance data (actual bandlimited RMS metric) is collected and J_{ol} (online optimization metric) is calculated. The short burst which is next is the pre-ID scan to maximize the ID SNR. It is followed by the four identification combs. During the next, long stretch of time, the OM is improving the controller, and finally, closed loop performance (actual bandlimited RMS metric) data is collected.

The FSAI code stores the first 60 controllers that are terminated before their performance is fully evaluated. The termination can be due to either controller saturation or because $J > J_{ol}$. In the run shown in Figure 4, of the first 60 controllers terminated, none were destabilizing when checked offline on an accurate model of the PGZt to PGRGZ loop. However, for the X Loop run, 11 of the first 60 were destabilizing. Considering the time history of the output for this run, shown in Figure 5, a number of destabilizing controllers were implemented. The online ID was not as accurate for this loop due to a smaller SNR during identification. Even with destabilizing controllers being implemented, a nearly 60% improvement was obtained and the SEM maintained the output magnitude at an acceptable level (it never saturates during controller improvement).

Table 2 shows the results (average for Z Loop+D1) for the different loops and disturbances. The numbers reported are for the 0.5 to 50 Hz bandlimited RMS, not the online performance metric of $\sum_{k=1}^{600} y_k^2$. It is noted that the FSAI architecture did not need to be changed for the different cases. All that was changed in going between the different disturbances and loops was that in going from the Z loop to the X loop, a different GMC controller was used to initialize the MSH algorithm. The improvements

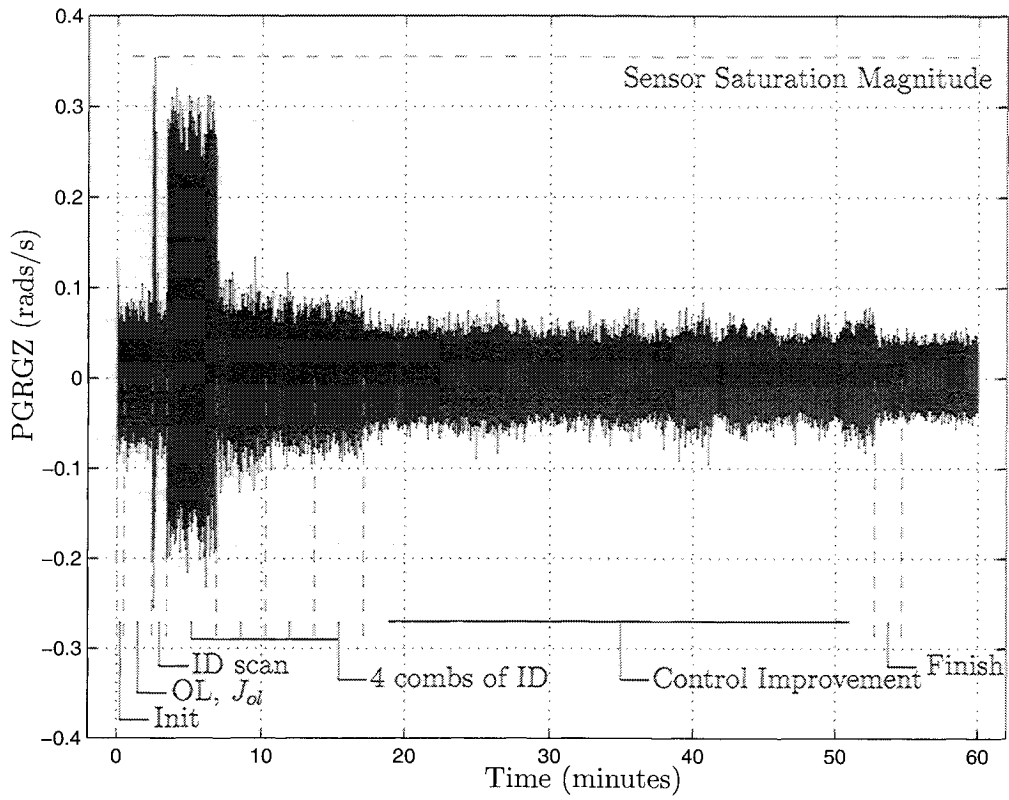


Figure 4: Output History for FSAI Control of Z Loop with Disturbance D2

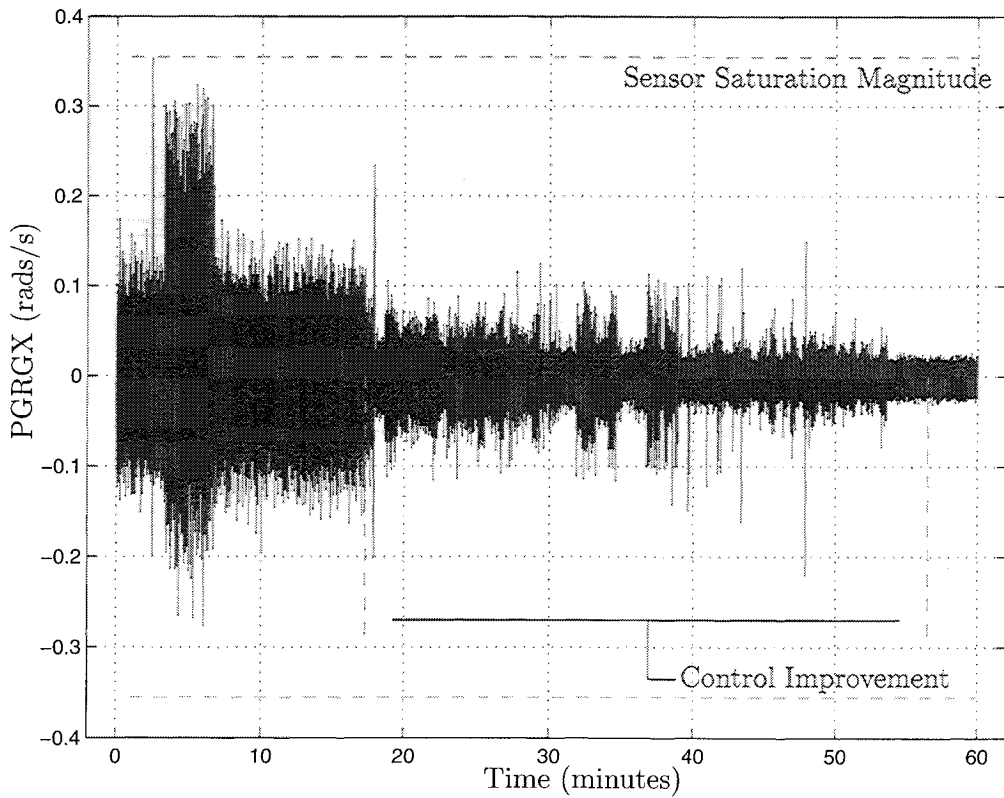


Figure 5: Output History for FSAI Control of X Loop with Disturbance D2

ranged from 23% to 58%.

Table 2: Summary of Results for X and Z Loops and Disturbances D1 and D2

| Controller | $100 \cdot J$ (rads ² /s ²) for: ^a | | |
|------------|--|------|-----------|
| | Z Loop/D1 | /D2 | X Loop/D2 |
| Open Loop | 3.01 | 2.41 | 4.23 |
| GMC | 1.82 | 1.72 | 1.72 |
| FSAI | 1.09 | 1.18 | 0.71 |
| % Improved | 40.1 | 31.4 | 58.7 |

^aSensor noise constrains $100 \cdot J > 0.15$

The PSD for the Z loop run with disturbance D1 is shown in Figure 6. The FSAI control algorithm has found a controller with worse low frequency performance, but this low frequency performance has been traded for better high frequency performance such that the overall performance is improved. This trade-off is observed in each of the runs and is probably due to the integral performance constraints of Freudenberg and Looze.⁶ Freudenberg and Looze derive integral constraints for non-minimum phase systems which imply that if the performance is increased in one frequency range, then it is necessarily decreased in another, and further, that the performance will be decreased by an amount which cannot be made arbitrarily small.

Conclusions

The FSAI approach provides a simple quasi-adaptive framework for online control improvement; it can be used to augment a low-to-medium gain, baseline robust controller. Further, the optimization algorithm and controller parameterization can be chosen independent of stability concerns. The version of FSAI for which theoretical guarantees have been derived has conservatively long settling times (based on a bound on the matrix exponential on a set of matrices), and ways for reducing this conservatism are currently under study (for instance, assuming the plant is always observable on Θ).

In practice, settling times were determined through trial and error, and an online, frequency-domain identification algorithm was developed to be used with an online Nyquist check. The Nyquist check provides a valuable pre-filter, but FSAI control can be made to work without it.

The experimental runs presented, while they do not characterize the performance of FSAI (this depending to a large part on the performance metric, optimization algorithm and controller parameterization), validate the overall approach: over 32,000 controllers were generated and selectively implemented in 3.4 hours of online control improvement without the output saturating. In fact, this behavior was achieved even though destabilizing controllers were

implemented. And in these runs baseline controllers were improved by 23 to 58%.

References

- [1] D.S. Bayard, F.Y. Hadaegh, Y. Yam, R.E. Schied, E. Mettler, and M.H. Milman, "Automated on-orbit frequency domain identification for large space structures," *Automatica*, vol. 27(6), pp. 931–946, 1991.
- [2] D.S. Bayard, Y. Yam, and E. Mettler, "A criterion for joint optimization of identification and robust control," *IEEE Trans. Automat. Contr.*, vol. 37(7), pp. 986–991, 1992.
- [3] M.E. Campbell, J.P. How, S.C.O. Grocott, and D.W. Miller, "On-orbit closed-loop control results for the middeck active control experiment," *J. Guid., Contr., and Dynam.*, vol. 22(2), pp. 267–277, 1999.
- [4] E.F. Crawley, M.S. Barlow, M.C. van Schoor, and B. Masters, "Measurement of the modal parameters of a space structure in zero gravity," *J. Guid., Contr., and Dynam.*, vol. 18(3), pp. 385–394, 1995.
- [5] G.F. Franklin, J.D. Powell, and M. Workman, *Digital Control of Dynamic Systems*. Menlo Park, CA: Addison Wesley Longman, Inc., third ed., 1998.
- [6] J.S. Freudenberg and D.P. Looze, "Right half plane poles and zeros and design tradeoffs in feedback systems," *IEEE Trans. Automat. Contr.*, vol. 30(6), pp. 555–565, 1985.
- [7] E. Gilbert and K.-T. Tan, "Linear systems with state and control constraints: The theory and application of maximal output admissible sets," *IEEE Trans. Automat. Contr.*, vol. 36(9), pp. 1008–1020, 1991.
- [8] S. Grocott, J. How, D. Miller, D. MacMartin, and K. Liu, "Robust control design and implementation on the Middeck Active Control Experiment," *J. Guid., Contr., and Dynam.*, vol. 17(6), pp. 1163–1170, 1994.
- [9] Y.-C. Ho, "The no free lunch theorem and the human-machine interface," *IEEE Contr. Syst. Magazine*, vol. 19(3), pp. 8–10, 1999.
- [10] R. Horst and P.M. Pardalos, eds., *Handbook of Global Optimization*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1995.
- [11] P.A. Ioannou and J. Sun, *Robust Adaptive Control*. Upper Saddle River, NJ: Prentice Hall, Inc., 1996.
- [12] R.D. Kaminsky and T.E. Djaferis, "The finite inclusions theorem," *IEEE Trans. Automat. Contr.*, vol. 40(3), pp. 549–551, 1995.




Figure 6: Open, GMC and Final FSAI PGRGZ PSDs for Z Loop with Disturbance D1

- [13] R.L. Kosut, M.K. Lau, and S.P. Boyd, "Set-membership identification of systems with parametric and nonparametric uncertainty," *IEEE Trans. Automat. Contr.*, vol. 37(7), pp. 929–941, 1992.
- [14] D. Liberzon and A.S. Morse, "Basic problems in stability and design of switched systems," *IEEE Contr. Syst. Magazine*, vol. 19(5), pp. 59–70, 1999.
- [15] D.W. Miller, E.F. Crawley, J.P. How, K. Liu, M.E. Campbell, S.C.O. Grocott, R.M. Glaese, T.D. Tuttle, G. Stover, J.A. Woods-Vedeler, J. de Luis, E. Bokhour, R. Grimes, K. Scholle, C. Krebs, and R. Renshaw, "The middeck active control experiment (MACE): Summary report," tech. rep., MIT SERC#7-96, 1996.
- [16] R.R. Ninneman and K.K. Denoyer, "Middeck active control experiment reflight (MACE II) reflight program," in *Proc. AIAA Space Tech. Conf.*, AIAA Paper#99-4584, Albuquerque, NM 1999.
- [17] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in C*. Cambridge, United Kingdom: Cambridge University Press, second ed., 1992.
- [18] D.P. Scharf, *Fixed-Structure Adaptive-Improvement Control: Performance Enhancing, Indirect Adaptive Control for Stable Plants*. PhD thesis, University of Michigan, 2001.
- [19] T.T. Tay and J.B. Moore, "Enhancement of fixed controllers via adaptive-Q disturbance estimate feedback," *Automatica*, vol. 27(1), pp. 39–53, 1991.
- [20] A. Törn and A. Žilinskas, *Global Optimization*. New York, NY: Springer-Verlag, 1989.
- [21] C. Van Loan, "A study of the matrix exponential," tech. rep., Univ. of Manchester Numerical Analysis Report 10, Manchester, England, 1975.
- [22] Z. Wang, I.M.Y. Mareels, and J.B. Moore, "Adaptive disturbance rejection," in *Proc. IEEE Conf. on Decision and Cont.*, (Brighton, England), pp. 2836–2841, 1991.
- [23] B. Wie and K.-W. Byun, "New generalized structural filtering concept for active vibration control synthesis," *J. Guid., Contr., and Dynam.*, vol. 12(2), pp. 147–154, 1989.
- [24] Y. Yam, D.S. Bayard, and R.E. Scheid, "Frequency domain identification for robust large space structure control design," in *Proc. American Cont. Conf.*, (Green Valley, AZ), pp. 3021–3023, 1991.
- [25] G. Zhai, B. Hu, K. Yasuda, and A.N. Michel, "Stability analysis of switched systems with stable and unstable subsystems: An average dwell time approach," in *Proc. American Cont. Conf.*, (Chicago, IL), pp. 200–204, 2000.